



**AAS 02-024**

## Descent-Speed Testing of a Hazard Detection System for Safe Landing on Mars

Andrew Edie Johnson and Eli David Skulsky

Jet Propulsion Laboratory

---

### 25th ANNUAL AAS GUIDANCE AND CONTROL CONFERENCE

---

February 6-10, 2002  
Breckenridge, Colorado

Sponsored by  
Rocky Mountain Section



AAS Publications Office, P.O. Box 28130 - San Diego, California 92198

## DESCENT-SPEED TESTING OF A HAZARD DETECTION SYSTEM FOR SAFE LANDING ON MARS

**Andrew Edie Johnson and Eli David Skulsky \***

A prototype hazard detection system under development at JPL for the next generation of Martian landers consists of a scanning laser radar (LIDAR), inertial measurement unit (IMU), and hazard detection and avoidance algorithms. The hazard detection algorithms construct a topographic map of the landing area using the motion-compensated LIDAR data and subsequently generate estimates of surface slope and roughness. These estimates are used to identify areas that exceed landing constraints dictated by vehicle design. The hazard detection software then selects a safe landing site and this location is provided to the guidance algorithms, which then steer the vehicle to the selected location.

A series of tests were performed to demonstrate the capabilities of the prototype hazard detection system at realistic terminal descent speeds. The test methodology was described in detail in a prior paper; in this paper we review the test methodology, describe the hazard detection algorithms, and provide preliminary results from the first hazard detection test.

The ultimate goal of the hazard detection tests is to verify the performance of a LIDAR-based system for Mars landing. The performance of the system is characterized through post-test data processing as follows. First, the flight sensor data was processed to determine the vehicle trajectory, the shape of the imaged target and the detected hazards. These "flight products" were then compared to the truth trajectory generated from truth sensor measurements, the truth shape of the target, and hazards known to exist in the target. Differences between the flight and ground products were then used to determine the overall system performance as well as any weak links in the data collection system. Preliminary results a shape reconstruction error less than 0.1m. Given correct trajectory data, the system was capable of detecting all hazards with radii greater than 0.5 m.

---

\* The authors are with the Jet Propulsion Laboratory, California Institute of Technology. All correspondence should be sent to Andrew Edie Johnson, Jet Propulsion Laboratory, 4800 Oak Grove Drive, M/S 125-209, Pasadena, CA 91109, e-mail: aej@jpl.nasa.gov.

## **INTRODUCTION**

In a previous paper [1] we described an approach for testing a prototype system designed to detect hazards during the terminal descent phase of a mission to Mars. The system under development at the Jet Propulsion Laboratory (JPL) utilizes a scanning laser radar and an inertial measurement unit to map the landing zone, identify hazards, and then guide the vehicle to a safe landing site. The test approach consisted of placing this system on a rocket sled, accelerating the sled to typical descent speeds, and mapping a simulated Martian terrain. The data would then be processed to identify hazards and locate a safe landing site.

Since ref. 1 was written several such tests have been performed. In this paper we provide an update on the testing methodology and we present preliminary results from one of the tests.

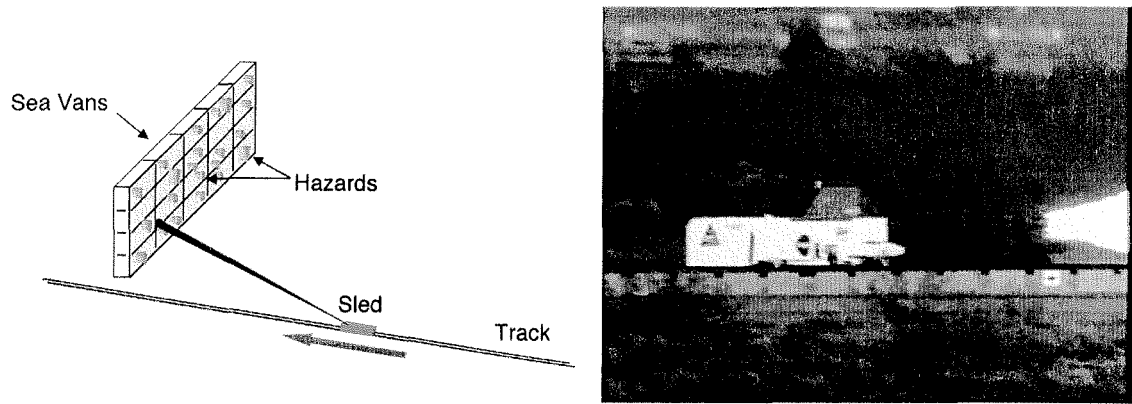
## **HAZARD DETECTION SYSTEM**

The prototype JPL hazard detection system consists of a scanning laser radar (LIDAR), inertial measurement unit (IMU), and hazard detection and avoidance algorithms. The LIDAR performs a raster scan of the landing zone and data from the IMU is used to compensate for rotational and translational motion during the scan. The hazard detection algorithms construct a topographic map of the area using the motion-corrected LIDAR data and generate estimates of surface slope and roughness. These estimates are subsequently used to identify areas that exceed landing constraints dictated by vehicle design and the hazard detection software then selects a safe landing site. For an actual Mars landing mission, this safe-landing location would be provided to the guidance algorithms which would steer the vehicle to the selected location.

## **Sled Test Overview**

To test the hazard detection algorithms with real sensor data collected at typical descent speeds, the hazard detection system was placed on a rocket sled. The sled was accelerated to approximately 60 m/sec and measurements were obtained beginning approximately 500 m in front of the target and continuing until the sled passed the target. The on-board sensors were mounted to an optical bench that was secured to the sled via a passive vibration isolation system designed and built at JPL.


Ground truth position and attitude information were obtained by independent sensor measurements from an onboard digital video camera, six onboard accelerometers, and a track-mounted laser range finder. All sensor data (flight-like and ground-truth measurements) were time-tagged and synchronized to better than a microsecond using a GPS time signal.



**Figure 1 LIDAR scanning as the sled travels along the track.**

The Martian terrain was represented by a set of hemispheres of various radii mounted to a wall of stacked sea vans (12 m  $\times$  65 m) placed on the side of the track. The hemispheres were mounted to plywood boards hung on the side of the sea vans; the boards were used to facilitate target construction and surveying and hemispheres were selected to facilitate post-test analysis. Using the surveyed positions and radii of each hemisphere, a ground truth CAD model of the target was generated; this model was used to verify the output of the surface reconstruction and hazard detection algorithms.

Terminal descent speed for future Mars lander missions is expected to be between 50 and 100 m/sec when the vehicle is several hundred meters above the planet surface, so the trajectory was designed such that the sled speed 500 m from the target was within that range.



**Figure 2 Target wall photo and CAD model rendering.**

## Flight-Like Sensors

The two “flight-like” sensors—the IMU and the LIDAR—functionally resemble the sensors that would be used in the flight system of a Mars lander mission. For the tests we used a single Litton LN-200 Inertial Measurement Unit and the Optech Laser Rangefinding Instrument (LRI).

The LN-200 IMU consists of three orthogonal fiber-optic gyros and three orthogonal silicon accelerometers in a package that is roughly the size of a coffee mug. The LN-200 is configured to provide 400 samples per second (angular velocity and acceleration counts) via an RS-485 serial interface.

The Optech LRI incorporates an infrared laser altimeter (1064 nm) and two scanning mirrors to capture three-dimensional shape data. The LRI has a pulse repetition frequency (PRF) of 8 KHz, a maximum field of view of  $10^\circ \times 10^\circ$  with a maximum resolution of  $500 \times 1000$  samples, a maximum range of 500 m, and 2 mrad beam divergence. For the sled tests, a  $5^\circ \times 10^\circ$  and  $50 \times 100$  sample scan pattern was employed; this scan pattern was selected to provide 10 to 15 scans of the target during each test. A dedicated computer collected range and angle measurements from the LRI while the same onboard computer that sampled the IMU and accelerometer collected timing data for the LRI. As with the ground truth data, all measurements were time tagged and synchronized to better than a microsecond using a GPS time signal.




Figure 3 CAD model of rocket sled and instrument pallet. The diagram on the right shows the digital video camera, LIDAR, IMU, and accelerometers mounted on the passive vibration isolation system.

## Truth Sensors

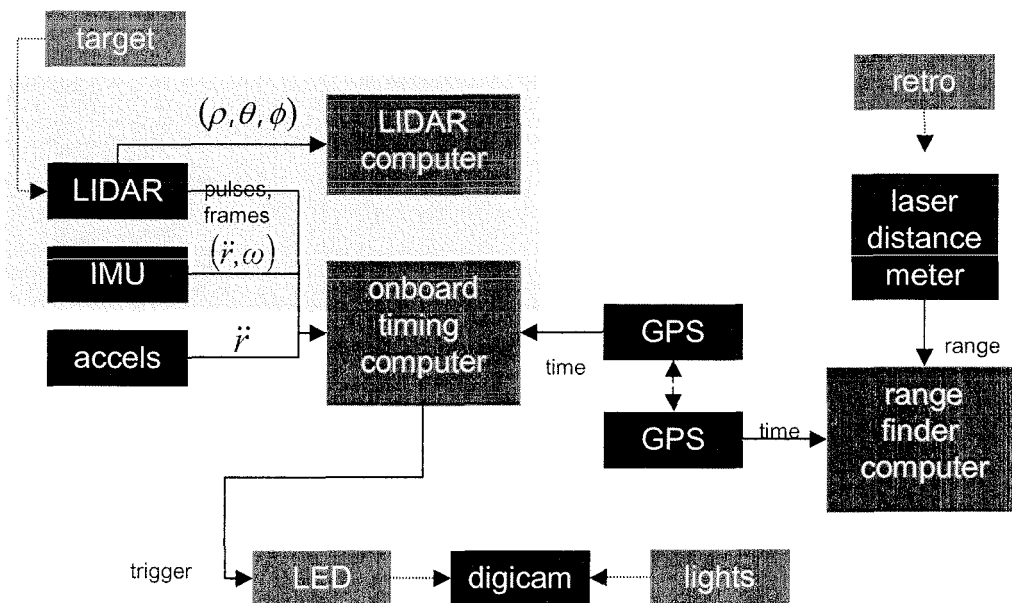
Measurements from the truth sensors are used to generate a high accuracy estimate of the sled trajectory that is completely independent of the flight sensor measurements. The truth sensors consist of a digital video camera (digicam), six QA-2000 Q-flex accelerometers, and a laser distance meter (LDM). The camera and the accelerometers were mounted on the optical bench; the LDM was fixed to the side of the track.

The digicam was used to obtain the sled attitude with respect to the target. The camera, which operates at 30 frames per second, captures images of a set of surveyed lights on and in front of the target. The camera images are internally time tagged and synchronized with other data products by illuminating an LED in front of the camera at a time controlled by the data collection system.

After each test the video images were analyzed to obtain the direction to the lights as seen by the camera. Next, the light direction measurements were combined with measurements of the camera position and knowledge of the three-dimensional position of the lights to obtain estimates of the sled-mounted optical bench versus time.

The QA-2000 Q-flex accelerometers serve two purposes: estimation of the sled position and estimation of the optical bench vibrational environment. Accelerometer data is collected and time tagged by the onboard field computer.

The laser distance meter (LDM), manufactured by Riegl USA, was used to provide direct (100 Hz) measurements of the sled position. The LDM measurements—which are accurate to about 2.5 cm—were collected by a computer and tagged with a GPS signal (the computer includes a GPS receiver). A retro-reflector on the front of the sled increases the operating range of the LDM to 700 m.



**Figure 4 Instrumentation block diagram.**

Figure 4 shows how the measurements from each sensor are time-tagged and stored with consistent GPS time for post-test analysis.

## Alignment

Since LIDAR frames are not captured instantly (each frame requires about one second to collect), motion correction must be performed on each pixel of each frame before the hazard detection software can be executed; consequently, the rigid transformation

between the LRI and the IMU must be known. Similarly, the orientation of the Q-flex accelerometers and high-speed camera must be known with respect to the LRI to validate the hazard detection results. Also required is the location (but not the orientation) of the laser distance meter and the target lights. To this end, an alignment procedure was developed to determine the orientation and position of any instrument with respect to any other instrument (see Figure 5). The objectives of the alignment procedure were to determine the position and orientation of all sensors in a Master Reference Frame at all times and to determine the orientation of the IMU and Q-flex accelerometers in geocentric inertial (GCI) coordinates to compensate for Earth rotational motion and gravity.

Along the track are permanent, surveyed, quarter-size “monuments” spaced roughly every meter. Because they are permanent and their locations are known, two of these monuments (along with a third permanent off-track marker) were used to construct a permanent coordinate frame to which every instrument and target was referenced. We call this coordinate frame the *Master Reference Frame* (MRF).

To obtain the transformation between an arbitrary coordinate frame (call it *A*) and the MRF, we use a total station surveying instrument (which provides the angles and range to a target) to measure the position in total station coordinates of the MRF monuments and the position in total station coordinates of the points which define *A*. From there it is straightforward to construct the transformation between the MRF and *A*. We did this for every instrument and for every target. Furthermore, we measured the position (though not the orientation because we treat these as points) of the laser distance meter and the target lights.




Figure 5 Test alignment diagram.

A similar procedure was performed for the instruments on the optical bench. In this case, we measured the orientation of each instrument the LIDAR, IMU, and accelerometers using a multiple theodolite system (MTS) in MTS coordinates; this step allows us to determine the orientation of any instrument with respect to the high-speed camera. Next, the high-speed camera is calibrated to correct for focal length and optical distortion. As described below, the attitude of the high-speed camera with respect to the target as the sled travels along the track can be determined by measuring the position of the target lights in the camera field of view. If the camera attitude is known and if the orientation of an instrument with respect to the camera is known, then we can determine the attitude of that instrument with respect to the target.

A final step in the alignment procedure is required to obtain the orientation of the IMU in geocentric inertial coordinates. This step is necessary to remove the biases introduced into IMU measurements by Earth rotation and gravity. Using a Differential GPS system and a total station this transformation was computed by measuring the position of distant monuments in GPS and total station coordinates followed by a correction that maps geodetic to geocentric coordinates.

## **POST-TEST DATA PROCESSING**

The ultimate goal of the rocket sled tests is to verify the performance of a LIDAR-based hazard detection system for Mars landing. Differences between the flight and ground products will be used to determine the overall system performance as well as any weak links in the data collection system. Figure 6 shows the block diagram for the three tests that will be used to assess system performance. The first test will be to determine if the 3-D shape generated from the flight sensors matches the ground truth shape. The second test will compare the results of hazard detection algorithms applied to the flight data and hazard detection applied to the ground truth shape. The final test will compare the trajectory generated by combining measurements from the flight IMU and LIDAR to the truth trajectory.

Temporary suspension of funding has prevented us from completing all of the tests on the data collected from the two rocket sled tests. We were only able to process data from the first rocket sled test (HDT1B), and, for this test, we were not able to generate the flight trajectory that incorporates IMU and LIDAR measurements. Instead we were only able to integrate the IMU to generate the flight trajectory. As will be shown below, the flight trajectory generated in this fashion is sufficient for estimating the attitude of the sled, but due to filter divergence, cannot estimate the position of the sled to sufficient accuracy for motion correction. Consequently, the tests we have performed have focused on comparing the shape and hazards detected using the trajectory where attitude comes from the digital camcorder and position comes from the laser distance meter (TRUTH-TRUTH) and a trajectory where attitude comes from the IMU and position comes from the LDM (FLIGHT-TRUTH). These tests will allow us to assess the system performance given the best possible trajectory generated from the truth sensors and system



performance when the attitude trajectory comes from the flight sensor. When funding continues for this effort, the remaining tests will be conducted.

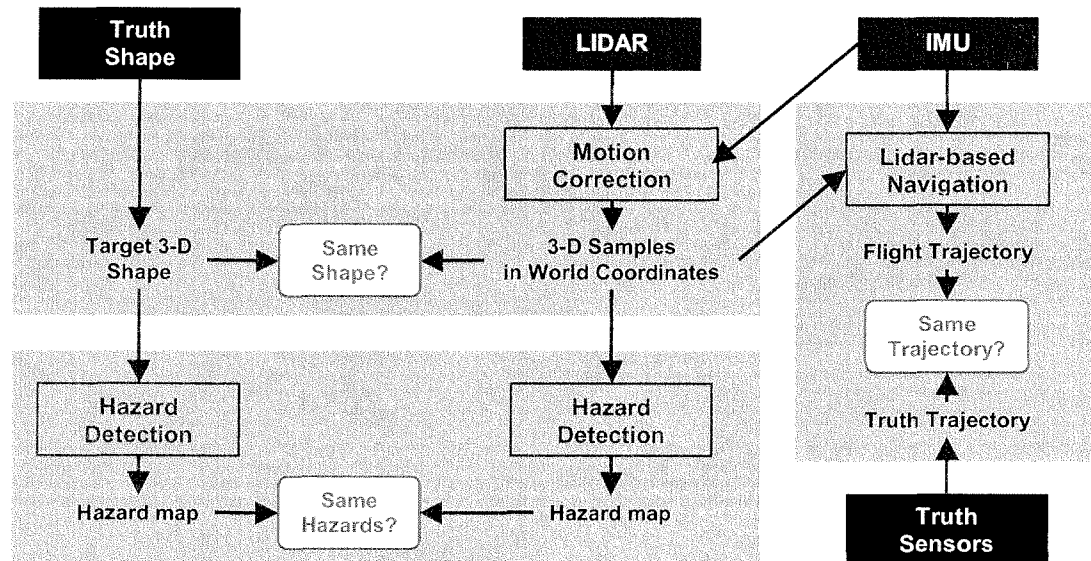


Figure 6 Post-test data processing block diagram

## Truth Trajectory Generation


Generating the truth trajectory occurs in two stages. First the position of the sled in the MRF is determined using the Laser Distance Meter (LDM). Next the attitude of the sled in the MRF is determined using the digital camcorder (digicam) images of the target lights. Together these measurements constitute the truth trajectory.

## Truth Position Generation

Sled position is estimated from the LDM measurements as follows (see Figure 7). The LDM is placed behind the target and close to the track. The LDM is pointed precisely down the track toward the sled start position so that range measurements can be obtained from the sled in the region of the track that is 500 m in front of the target. Retro-reflective prisms are placed on the front of the sled to obtain maximum range from the sensor and guarantee that the front of the sled is measured (and not other surfaces on the sled).

As stated previously, markers along the track are surveyed into the MRF prior to the test; these measurements define a piece-wise linear representation of the track. While the sled is on the track, the position of the sled retro-reflector, the optical bench and a few track markers are surveyed. Using this information, the transformation between the track markers and the sled retro-reflectors and subsequently the piecewise linear curve that the sled retro reflectors follow is computed. These measurements are also used to compute the transformation between the optical bench and the retro-reflector.

The LDM provides a measurement of the range to the sled retro-reflector along the line of sight; this measurement defines a sphere. Because the range center of the LDM is known in the MRF, this sphere is also defined in the MRF. By computing the intersection of this sphere and the piece-wise linear retro-reflector curve defined above, the position of the sled retro-reflector is determined for each LDM measurement (100 Hz). Finally, the transformation between the retro-reflector and the optical bench and the transformations between sensors on the optical bench are used to determine the position of each sensor (LRI, digicam, IMU) in the MRF for each LDM measurement. The LDM measurements and the onboard sensor measurements (IMU/LIDAR/digicam) are synchronized in time because they are all time-tagged using GPS.



**Figure 7 Procedure for computing sled position from track survey and laser distance meter (LDM) measurements.**


Figure 9 shows the truth position trajectory components computed from the LDM data collected during HDT1B.

### Truth Attitude Generation

The truth sled attitude is computed for every image taken with the digicam using an image-based attitude estimation algorithm. The inputs to the algorithm are the known 3-D positions of target lights in the MRF, the 2-D image positions of each light in the image taken by the digicam and the position of the camera computed using the LDM. The output of the estimation procedure is the attitude of the sensors on the optical bench, in the MRF, for each digicam image (30 Hz). The algorithm has three stages: image timing, light position extraction and attitude estimation.

Before each sled test, the digicam field of view is set to approximately  $12^\circ$  and the camera is focused at infinity. A standard calibration procedure [ref] is then employed to determine the mapping between image pixels and 3-D rays in the sled coordinate system. This procedure solves for camera internal parameters (field of view optical center, image skew and radial distortion) as well as position and attitude of the camera in the sled coordinate frame.

After the test, the sequence of imagery between LED activation and sled stop is digitized and stored as a sequence of separate images. The image that records the first light from the LED's (the LED image) is given the time recorded for LED activation by the onboard timing computer. The times of subsequent images are then computed based on the image index after the LED image and the 30 frames-per-second rate of the digicam.



**Figure 8 Digicam images used to estimate sled attitude. (top left) The LED image with lights in background. (top right) Example light input image. (bottom left) Connected components extracted from light image. (bottom right) Labeled centroids of lights.**

Once timed, each image is processed to extract the light positions in the image. First, a threshold (254 dN for HDT1B) is applied to the image to create a binary image of light and not light regions. An original image and the binary image resulting from a threshold is shown in Figure 8. Next, connected components, an algorithm that groups together adjacent pixels, is applied to the binary image. The resulting regions are shown in Figure 8. For various reasons (e.g., overlapping lights, stray light) some of the extracted regions may not correspond to a single light. These extraneous regions should be eliminated prior to pose estimation. Since lights are expected to be roughly circular, most of these regions can be eliminated by employing a test on connected component aspect ratio; if the ratio of the width and height of the connected component is less than a threshold (0.75) or greater than a threshold (1.25) it is eliminated from consideration. The final step of light extraction computes the centroids of each remaining connected component and labels

these regions in raster scan order. An example of labeled connected components is given in Figure 8.

In order to estimate the attitude of the camera, a correspondence between 2D light image centroids  $c_i$  and 3D light positions  $P_j$  must be established. Establishing this correspondence is nontrivial because lights can overlap in the image, the camera position and attitude change significantly between frames and lights appear and disappear during the test. Our approach to establishing the correspondences  $m(i) = j$  is to first project each  $P_j$  into the image using the current best estimate for the position and attitude of the camera; this results in some estimated 2D coordinates  $p_j$  for each light in the image. Next for each  $c_i$ , the closest  $p_j$ , based on 2D Euclidean distance, is computed to establish the correspondence mapping  $m(i) = j$ . In other words, image centroids are matched to the light that is closest when projected into the image.

Once the correspondence between centroids and lights is established, the attitude quaternion  $q$  of the camera is estimated by minimizing

$$\min_q \sum_i \|c_i - p(P_{m(i)}, q, T)\|^2$$

where  $P_{m(i)}$  is the 3D light position matched to the centroid  $c_i$ ,  $T$  is the position of the camera in the MRF given by the LDM, and  $p$  is the function that maps 3D points to 2D image coordinates. This function is minimized using the Levenberg-Marquardt non-linear minimization algorithm in a similar fashion as that described in [Szeliski and Kang].

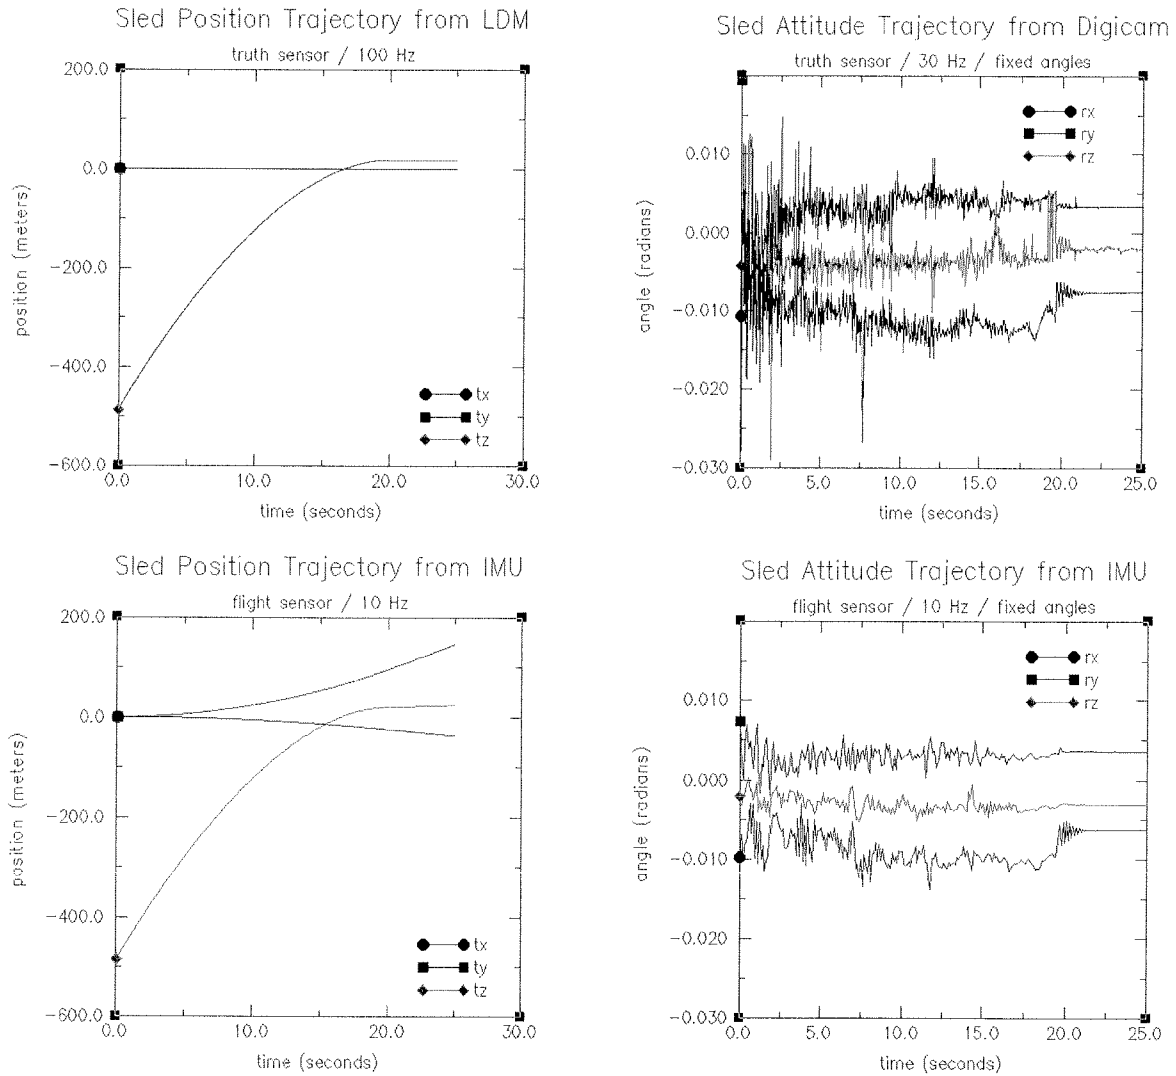
The above minimization will produce erroneous results if the initial correspondences are wrong. To mitigate this problem we place light correspondence and attitude estimation in a loop. First the correspondence between lights and centroids are established using the sensor position and the resting sensor attitude. These correspondences are used to estimate the new sensor attitude. Using this attitude, correspondences are reestablished. If no changes in correspondence occur the iterations are finished. Otherwise the attitude is re-estimated and the loop repeats. This type of algorithm is derived from the Iterative Closest Point algorithm that is commonly used for 3-D image registration [Besl and McKay].

The truth attitude trajectory computed for HDT1B is shown in Figure 9. The 6 degree-of-freedom truth trajectory used for motion correction is the combination of the 100 Hz position estimates from the LDM and the 30Hz attitude estimates from the digicam.

## Flight Trajectory Generation

Integrating the IMU data generates the flight trajectory. Using a model of the LN-200 IMU, the 3-axis acceleration and the 3-axis angular velocity are integrated simultaneously by a navigation filter to generate the position and attitude of the sled as a function of time. By combining this trajectory with sensor coordinate transformations and the initial starting position and attitude of the IMU (from ground truth measurement),

the position and attitude of the sled is transformed into the MRF. The filtered IMU position and IMU attitude trajectories are given in Figure 9.



**Figure 9** Computed truth (LDM and digicam) and flight (IMU) trajectories.

### Same Trajectory Test

Qualitative analysis of the attitude trajectories shown in Figure 9 indicates that the IMU derived attitude trajectory and the digicam-derived attitude agree quite well. The digicam attitude is generated at 30 Hz and no dynamic model is used to smooth the trajectory. Consequently, the digicam attitude contains more high frequency components than the 10 Hz IMU attitude trajectory. Some of the attitude spikes in the digicam trajectory may be due to incorrect estimation of the sled attitude due to incorrect correspondence between lights and centroids. Future analysis will indicate if this is the case.

The MRF is aligned with the track and the sled travels in the positive  $z$  direction. Therefore, the expected trajectory will be  $x$  and  $y$  components close to zero and the  $z$  component increasing. The position trajectory generated by the LDM follows this trend exactly while the position trajectory generated by the IMU does not. The IMU trajectory diverges as the sled approaches zero velocity. This is to be expected; near zero velocity, unmodeled noise sources and biases will dominate and the filter will diverge. Future work will mitigate this problem by incorporating velocity measurements from the alignment of consecutive LIDAR scans into the navigation filter.

Since the IMU derived position trajectory is useless for motion correction of the LIDAR samples, it will not be used in the remaining same shape and hazard detection analysis. Instead two trajectories will be used to perform the analysis: the trajectory using digicam attitude and LDM position (TRUTH-TRUTH) and the trajectory using IMU attitude and LDM position (FLIGHT-TRUTH). These tests will allow us to assess the system performance given the best possible trajectory generated from the truth sensors and system performance when the attitude trajectory comes from the flight sensor.


## **Truth Shape Generation**

The ground truth shape of the target is generated through a combination of surveying and construction of the target to tolerance. The target is composed of multiple plywood boards onto which are attached acrylic hemispheres. The hemispheres are constructed with a known tolerance on hemisphere diameter. The hemispheres are then placed accurately on the each plywood board within a known tolerance to the corners of the board. After the boards are hung on the wall of sea vans, the position and attitude of each board is determined by surveying the corners of each board. Using these surveying measurements and parametric models of the hemispheres, a CAD model of the entire target surface is constructed. A rendering of the CAD model for the target is given in Figure 2.

## **Motion Correction**

The LIDAR samples are timed and their 3D positions are known in the LIDAR reference frame. Figure 10 shows the LIDAR samples collected for HDT1B in the LIDAR reference frame. To be useful for hazard detection, the LIDAR samples must be placed in a surface fixed coordinate system by removing the motion that occurs during and between each scan. This process is called motion correction and for the sled tests the appropriate surface fixed coordinate system is the MRF.


If the trajectory of the LIDAR is known, then motion correction is simple. The time of each sample is used to interpolate between measurements in the trajectory to get a position and attitude of the LIDAR in the MRF for each sample. The resulting transformation is then applied to the 3D sample to place it in the MRF. Figure 10 shows the resulting motion correction of the LIDAR samples using the TRUTH-TRUTH and FLIGHT-TRUTH trajectories.



**Figure 10 Motion correction of LIDAR scans.** (left) Uncorrected scans shown in the LIDAR reference frame. (right) After motion correction using position and attitude trajectories, the LIDAR scans are aligned in the MRF.

### Same Shape Test

The next test of system performance is to compare the motion corrected LIDAR samples to the target CAD model. Figure 11 gives a qualitative feeling for how well motion correction performs. The TRUTH-TRUTH motion corrected samples appear to do a better job than the FLIGHT-TRUTH at aligning the samples with the target.



**Figure 11 Same shape comparison.** (top) 3D samples motion corrected using the LDM and digicam compared to the CAD target. (bottom) 3D samples motion corrected using the LDM and the IMU attitude compared to the CAD target.

To obtain a quantitative assessment of performance, the minimum Euclidean distance between each sample and the target is computed. The sea-vans that support the target were not surveyed, but, because many of the samples returned measurements from the

sea-vans and not the target boards, the sea-vans need to be added to the target for correct distance measurements. To this end, a plane is fit to the sphere centers of all of the hemispherical hazards. A plane is a good approximation since the target wall is relatively flat.

The distance between each sample and the target is then the minimum distance between the sample and the combination of the target plane and all of the hemispherical hazards. Suppose the plane has normal  $\mathbf{n}$ , oriented toward the MRF origin, and constant  $d$ . The distance  $r$  between the plane and sample  $s$  is  $r = \mathbf{n} \cdot \mathbf{s} + d$ . This distance will be positive for samples in front of the target and negative for those behind the target. This distance between the sample  $s$  and each hemispherical target with radius  $R_i$  and center  $\mathbf{c}_i$  is  $r = \|\mathbf{s} - \mathbf{c}_i\| - R_i$ . This distance will be positive for samples outside the hazards and negative for those inside. Consequently, the minimum distance between a sample and the target will be

$$r = \min(\mathbf{n} \cdot \mathbf{s} + d, \min_i (\|\mathbf{s} - \mathbf{c}_i\| - R_i))$$

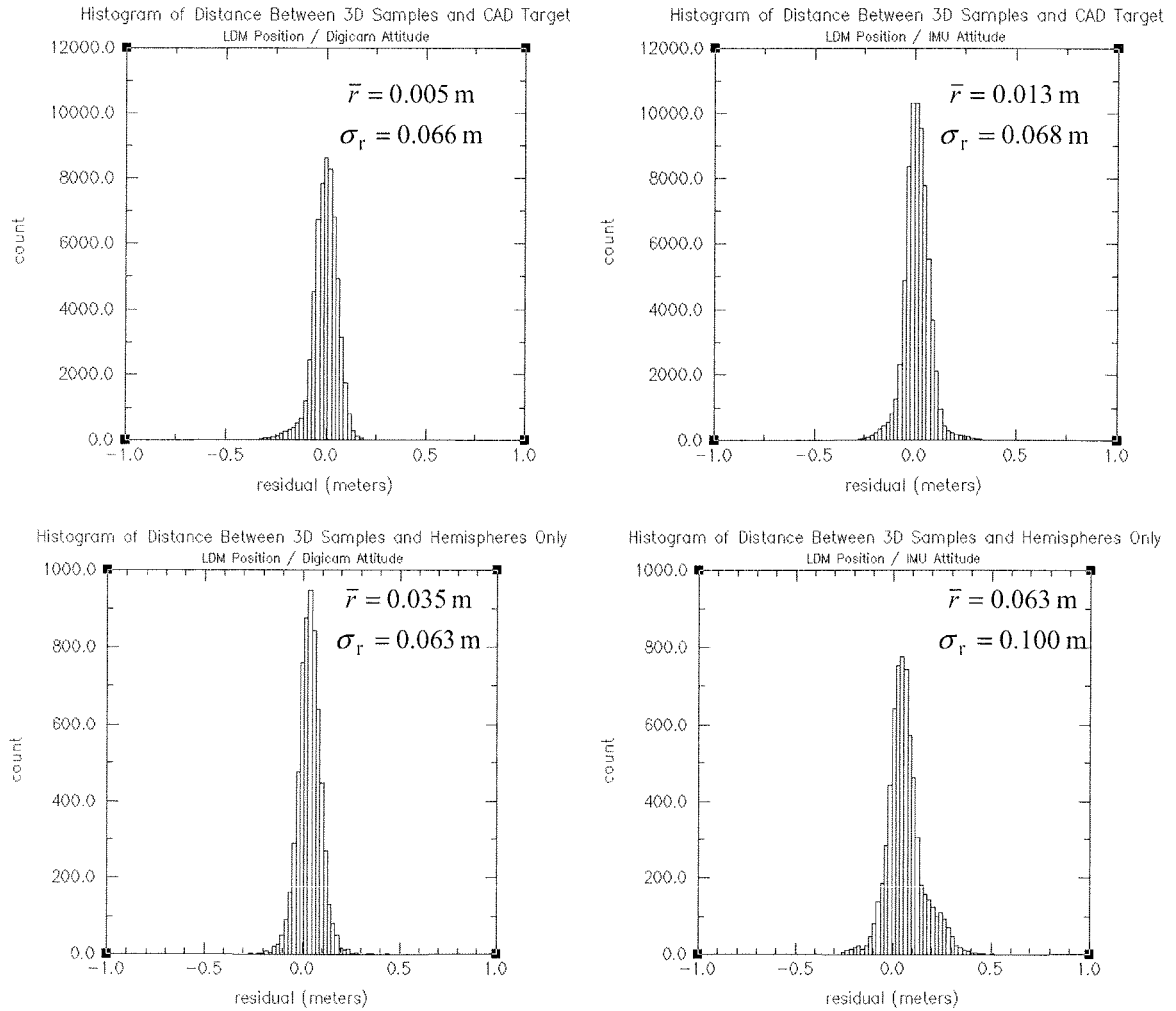
Histograms, means and standard deviations of the distance residuals for the TRUTH-TRUTH and FLIGHT-TRUTH motion corrected samples are given in Figure 12. Two forms of the histograms are given. The first contains residuals for all samples; this histogram shows that the total standard deviation of the residuals is small ( $\sim 0.06\text{m}$ ) given the size (0.5m diameter or greater) of the hazards that must be detected for safe landing. The second histogram shows the residuals for just those samples that are closer to the hazards than the target plane. The purpose of these histograms is to assess the performance of the LIDAR near the hazards that must be detected. The TRUTH-TRUTH histogram is still quite tight (0.063m standard deviation). However, the FLIGHT-TRUTH histogram is broader with a standard deviation of 0.1m. This indicates that the motion correction is not perfect and as a result, the samples are misplaced from their true positions. These errors are large enough that LIDAR-based navigation seems to be required for accurate motion correction and robust hazard detection.

## Hazard Detection

The final test of system performance compares the hazards detected from LIDAR samples to the true hazards as indicated by the target CAD model. A particular patch of terrain presents a hazard to the spacecraft during landing if the slope of the patch is too steep or the patch contains rocks or other protuberances that are taller than a certain terrain height. To quantitatively determine if a patch is hazardous, the slope and terrain variation over the patch must be measured. The hazard detection and avoidance algorithms used in this test are described in detail in [6]. In brief, our algorithms estimate the location of surface hazards given the motion-compensated flight shape as follows: first an elevation map is generated from the 3-D LIDAR samples. Next, estimates of local slope and roughness are computed over the entire elevation map. Finally, areas of the terrain map that exceed constraints on surface slope and roughness given the footprint of the lander are determined. Images of terrain, local slope, local roughness and detected



hazards for the TRUTH-TRUTH and FLIGHT-TRUTH motion corrected samples are given in Figure 13.



**Figure 12 Histogram of 3D distance residuals between motion corrected scanned data and target CAD model. Top histograms show all residuals while bottom histograms show residuals only for samples that are closer to a hemisphere than the background target plane.**

The terrain elevation maps are coded according to the color spectrum so that tall terrain features are red and low terrain features are blue. The three large 1m diameter hazards clearly show up in both maps. Because of the sled test geometry, the number of samples per unit area increases from right to left due to multiple scan overlap. Consequently averaging of samples in each terrain map bin will make the elevation data more accurate on the left of the terrain maps. This effect can be seen clearly in the terrain maps; on the left individual target boards can be clearly made out while on the right they become indistinguishable from the underlying terrain.

In the roughness (slope) map, roughness (slope) is colored from green to red. Red corresponds to regions that exceed the roughness (slope) constraint of 0.5 m (15°) or

regions where surface parameters cannot be computed due to lack of data. Yellow regions correspond to somewhat rough (sloped) regions and green corresponds to regions where the roughness (slope) is close to zero. Both roughness maps clearly show all three 1m hazards and most of the 0.5 m hazards. The slope of the target wall is much less than the  $15^\circ$  constraint, so both slope maps are entirely green.

The slope and roughness maps are combined using a weighted average to create a cost map. If either of the slope or roughness values computed in a terrain map bin exceeds its hazard tolerance threshold, all bins in the cost map within the size of the lander of this bin are given a value 2.0. Otherwise the cost map bins are set to the sum of the roughness normalized by the roughness threshold and the slope normalized by the slope threshold. The safe landing maps show the regions of the terrain map that are safe (green), unsafe (red) and unknown (yellow) according to the cost map.

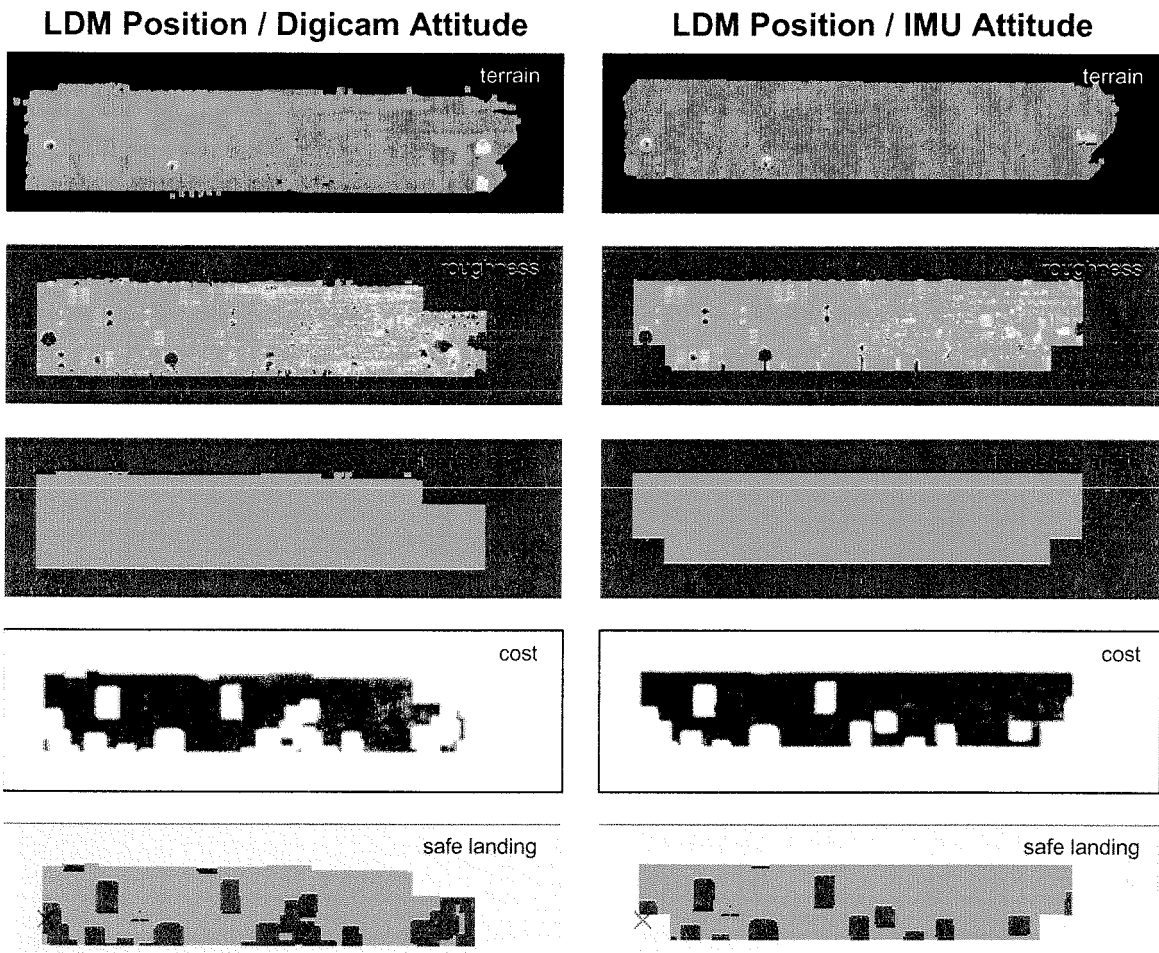


Figure 13 Hazard Detection parameter maps and results for two trajectories.

## Same Hazards Test

After hazards are detected in the terrain maps generated from LIDAR samples, they can be compared to the hazards that actually exist in the terrain map. Figure 14 show safe landing maps texture mapped onto the elevation maps for the TRUTH-TRUTH and FLIGHT-TRUTH LIDAR samples. A wireframe target CAD model is also superimposed for comparison. In the TRUTH-TRUTH data set all 20 of that hazards with radii greater than the hazard tolerance threshold (0.5 m) are correctly identified. In the FLIGHT-TRUTH data set all but one of the hazards are correctly identified. The missed hazard has a radius of 0.5 m, and Figure 14 it is in the distant portion of the terrain map where the data quality is poor due to sparse sampling of the target.

Given these results, LIDAR-based hazard detection appears quite promising. Given accurate trajectory data, it can reliably detect hazards that are large enough to harm a lander. It can do this while traveling at distances and descent speeds expected for the next generation of Mars landers.




Figure 14 Safe landing maps texture mapped onto terrain maps shown with wireframe rendering of target CAD model.

## CONCLUSION

An approach to testing a hazard detection system for landing on Mars was described. The approach provides a safe and repeatable means to demonstrate that LIDAR-based

hazard detection—a critical technology component of the next generation of Mars landers—is viable

Future work in this area will include modifications to the LIDAR system to incorporate an external mirror for landing site tracking during landing as well as tests using different LIDARs and alternate sensors including a radar.

## **ACKNOWLEDGMENTS**

This work was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

## **REFERENCES**

- [1] Skulsky, E.D., Johnson A.E., et al, "Rocket Sled Testing of a Prototype Terrain-Relative Navigation System," AAS 01-026, 24<sup>th</sup> Annual AAS Guidance and Control Conference, January 31 – February 4, 2001.
- [2] Golombek, M.P., Cook, R.A., et al, "Overview of the Mars Pathfinder Mission and Assessment of Landing Site Predictions," *Science*, December 5, 1997; 278: 1743–1748.
- [3] M.C. Malin and K.S. Edgett, "Sedimentary Rocks of Early Mars," *Science* 2000 December 8; 290:1927-1937.
- [4] M.C. Malin and K.S. Edgett, "Evidence for Recent Groundwater Seepage and Surface Runoff on Mars," *Science* 2000 June 30; 288: 2330-2335.
- [5] Johnson, A E., and San Martin, M., "Motion Estimation from Laser Ranging for Autonomous Comet Landing," *Proc. Int'l Conf. Robotics and Automation (ICRA '00)*, April 2000.
- [6] Johnson, A.E., Klumpp, A., Collier, J., and Wolf, A., "Lidar-Based Hazard Avoidance for Safe Landing on Mars," AAS/AIAA Space Flight Mechanic Meeting, (To Appear). February 2001.